

Quantifying the Impact of Parameter Tuning on Nature-Inspired Algorithms

Matthew Crossley, Andy Nisbet and Martyn Amos

School of Computing, Mathematics and Digital Technology,
Manchester Metropolitan University, Manchester M1 5GD, UK.
Email: m.crossley@mmu.ac.uk

Abstract

The problem of *parameterization* is often central to the effective deployment of nature-inspired algorithms. However, finding the optimal set of parameter values for a combination of problem instance and solution method is highly challenging, and few concrete guidelines exist on how and when such *tuning* may be performed. Previous work tends to either focus on a specific algorithm or use benchmark problems, and both of these restrictions limit the applicability of any findings. Here, we examine a *number* of different algorithms, and study them in a “problem agnostic” fashion (i.e., one that is not tied to specific instances) by considering their performance on *fitness landscapes* with varying characteristics. Using this approach, we make a number of observations on which algorithms may (or may not) benefit from tuning, and in which specific circumstances.

Introduction and Background

There exist many algorithms that are inspired by nature, and each has associated with it a set of *parameters*. These define specific features or details of an algorithm that may be altered in order to change the behaviour or performance of the method (for example, in evolutionary algorithms, parameters may include mutation rate or crossover probability). The problem of finding the optimal settings for these parameters (often referred to as “tuning”) is well-established (Lobo et al., 2007; Nannen et al., 2008; Akay and Karaboga, 2009; Birattari, 2009; Eiben and Smit, 2011), but little in-depth work has been performed on quantifying the *benefits* of tuning for a *range* of algorithms. We address this in the current paper, by investigating the precise benefits (or otherwise) of tuning for a number of different algorithms. Moreover, we do this in a way that is independent of any specific *problem*, by using an approach based on fitness landscape characteristics.

Previous work (Crossley et al., 2013) characterised a number of nature-inspired algorithms according to their performance on fitness landscapes with different features. However, the authors used the *default* parameter settings for each algorithm, which fails to reflect the fact that, in practice, methods are usually *tuned* prior to serious use (Leung

et al., 2003; Adenso-Diaz and Laguna, 2006; Koster and Beney, 2007; Ridge and Kudenko, 2010). Here, we extend this work by quantifying the relative merits of tuning for a range of algorithms in a wide variety of fitness landscape scenarios. We achieve this by assessing both their tuned and untuned behaviour, using the methods described in Crossley et al. (2013).

In order to tune our selected algorithms, we use the notion of *racing*, which was first introduced in the field of machine learning (Maron and Moore, 1993, 1997). Specifically, we use the F-race algorithm (Birattari et al., 2002; Yuan and Gallagher, 2004; Smit and Eiben, 2009; Birattari et al., 2010), which has been extensively used to find the best possible set of parameter values for a given problem in a limited time.

The rest of the paper is organised as follows: we first describe our approach in the Methodology section, before presenting our experimental findings in the Results section. We conclude with a discussion of the implications of our results, and suggest further work.

Methodology

Our methodology may be summarised as follows: (1) select a number of nature-inspired algorithms, and obtain consistent source code for their implementation; (2) for each algorithm, find the best parameter settings (i.e., tune) over a number of different problems; (3) compare the performance of tuned and untuned algorithms.

Algorithm selection

We compare a number of nature-inspired algorithms, all of which are commonly applied to continuous function optimisation (we use the same set as in Crossley et al. (2013)). These may be classified (Brabazon and O’Neill, 2006) as either *social*, *evolutionary* or *physical*. The social algorithms we select are Bacterial Foraging Optimisation Algorithm (BFOA) (Passino, 2002), Bees Algorithm (BA) (Pham et al., 2006), and Particle Swarm Optimisation (PSO) (Kennedy and Eberhart, 1995). The evolutionary algorithms selected are Genetic Algorithms (GA) (Goldberg, 1989) and Evolu-

tion Strategies (ES) (Bäck and Schwefel, 1993), and physical algorithms are represented by Harmony Search (HS) (Geem and Kim, 2001). We also include stochastic hill climbing (SHC) as a “baseline” algorithm; in contrast to Crossley et al. (2013) we exclude random search, as it has no parameters to tune. As before, we heed the observation that “Ideally, competing algorithms would be coded by the same expert programmer and run on the same test problems on the same computer configuration” (Barr et al., 1995). With that in mind, we use only implementations provided by Brownlee (2011). Space prevents a complete description of specific implementation details for each algorithm, but full implementation details can be found in Brownlee (2011), which is freely available and contains complete source code.

Tuning

Our fundamental goal is to investigate the pre- and post-tuned performance of our selected algorithms on landscapes with different general features, and thus identify characteristics of landscapes for which tuning may yield significant differences in algorithm performance. As Morgan and Gallagher (2010) observe, “Different problem types have their own characteristics, however it is usually the case that complementary insights into algorithm behaviour result from conducting larger experimental studies using a variety of different problem *types*” (our emphasis). Rather than using arbitrary benchmark instances of problems in order to perform tuning, we use a landscape-based approach, as utilised in Crossley et al. (2013). As Morgan and Gallagher (2010) explain, this Max-Set of Gaussians (MSG) method (Gallagher and Yuan, 2006) is a “randomised landscape generator that specifies test problems as a weighted sum of Gaussian functions. By specifying the number of Gaussians and the mean and covariance parameters for each component, a variety of test landscape instances can be generated. The topological properties of the landscapes are intuitively related to (and vary smoothly with) the parameters of the generator.” We now describe the characteristics under study:

Ruggedness of a landscape is often linked to its difficulty (Jones and Forrest, 1995), and factors affecting this include (1) the *number* of local optima (Horn and Goldberg, 1994), and (2) *ratio* of the fitness value of local optima to the global optimal value (Malan and Engelbrecht, 2009; Merz, 2000). Other significant factors concern (3) *dimensionality* (Hendtlass, 2009) (that is, the number of variables in the objective function), (4) *boundary constraints* (that is, the limits imposed on the value of a variable) (Kukkonen and Lampinen, 2005), and (5) *smoothness* of each Gaussian curve (effectively the gradient) used to generate the landscape (Beyer and Schwefel, 2002) - a smaller value indicates a smoother gradient. For each characteristic, we use the same ranges as in Crossley et al. (2013), summarised in Table 1.

To produce a test set of problems, we use the MSG landscape generator. For every value of every characteristic (in

Table 1: A summary of the ranges selected for the various characteristics in the landscape generation methodology.

Characteristic	Min	Max	Step	Default
No. of local optima	0	9	1	3
Ratio of local optima to global optimum	0.1	0.9	0.2	0.5
Dimensionality	1	10	1	2
Boundary constraints	10	100	10	30
Smoothness	10	100	10	15

the range specified in Table 1) we generate a set of five landscapes, which makes up the initial problem set for each value. We then use the F-racing methodology (Birattari et al., 2002) to find optimised parameters for each algorithm, over *every* value of *every* landscape characteristic used. We ensure that termination criteria are standardised, in order to ensure reasonable comparisons, and therefore use the number of objective function evaluations to determine when to terminate an algorithm’s run. We established, through initial experiments, that all selected algorithms generally converge within 20,000 objective function evaluations, so we use that as the specific value.

Comparison

We run each algorithm 100 times on each landscape in the set of landscapes generated for each particular characteristic value (when investigating smoothness, for example, we generate 1000 different landscapes (100 each for smoothness = 10 ... 100), and run each algorithm 100 times on each landscape). This is done first for all algorithms with ‘default’ parameter configurations, and then again, this time using the parameter configurations obtained through the F-Racing process. We measure the performance of each algorithm in terms of the mean (μ) and standard deviation (σ) of the exact average error obtained, over all values for a particular characteristic. That is, we investigate the *robustness* of each algorithm to changes in the values for each characteristic, rather than their absolute performance on specific problem instances. This allows us to identify specific landscape features where tuning may make a significant difference, some difference, or no difference at all, for a particular algorithm.

Results

We find that the effect of tuning is varied across algorithms, and that they fit into three categories: Algorithms which *do not* benefit from tuning (ES), algorithms which only benefit significantly from tuning when a landscape is ‘difficult’ for the algorithm using default parameters (BA, HS, PSO), and algorithms which *always* benefit from tuning (BFOA, GA, SHC). We summarise our results in Table 2; the full datasets

are available online¹. We now examine in detail the performance of each algorithm; for each graphical spider plot, the further a line is from the origin, the *smaller* the average error (that is, the “larger” an area, the larger the degree of robustness, which is considered “better”).

Bacterial Foraging Optimisation Algorithm

There exists little discussion on the role of different parameters in the BFOA. While some elements of the search pattern are clearly altered by various parameters, it is very difficult to estimate values for these. In the original description of the BFOA (Passino, 2002), the parameter values were assigned based on observation of actual bacterial colonies. While this may be true to the nature-inspired concept, it is not necessarily the best way to obtain optimal performance from the algorithm. The combination of parameters offered by BFOA gives a highly configurable search environment. Parameters such as step size and population size directly affect the potential area the algorithm can explore in a given number of objective function calculations. Attraction and repulsion weights, and the “space” over which these attraction and repulsion effects spread, work to control local optima avoidance. Parameters controlling the number of chemotactic steps before a reproduction step, and the number of reproduction steps before an elimination-dispersal event, control the balance of *exploitation* versus *exploration*. Given that the search behaviour of the algorithm is highly configurable, it is unsurprising that BFOA is heavily reliant on tuning.

Results for BFOA are shown in Figure 1. Across all characteristics, tuning offers a *significant* improvement on the average error and standard deviation of the performance - in many cases, improving the ranking of the algorithm from the largest average error to one of the smallest, and coping well with the changing characteristics. We see the most significant improvement where boundary constraint ranges change, a characteristic that is heavily reliant on parameters which control the range at which new solutions are generated (in the case of BFOA, this is the *step size*). Improvements are also shown for dimensionality and smoothness coefficient, increasing the performance of BFOA where there is little gradient information in a large fitness landscape. Smaller improvements are demonstrated by the increasing number of local optima and the increasing attractiveness of these local optima, but tuning still benefits the algorithm considerably.

In terms of the configurations selected by F-Racing, there is little variation in parameter values as characteristics change. Across all characteristics, and all values for those characteristics, there are only eight different configurations selected by racing. This suggests that, while it is difficult to find a *good* configuration, once it has been found, it is likely to be good for all *similar* problems. Tuning is vital to the

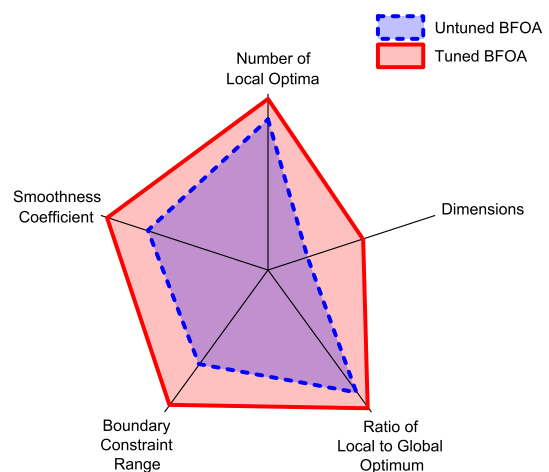


Figure 1: Summary of results for **bacterial foraging optimisation algorithm**.

performance of the BFOA, but it is possible that by exploring problems using a similar methodology to that demonstrated here, we may create a ‘bank’ of promising configurations.

Bees Algorithm

The BA is considered to be an algorithm on which parameterisation has little effect (Pham et al., 2006). We observe that the BA is one of the best *untuned* performers in this study, offering weight to this argument for relative parameter insensitivity. In terms of adjusting the BA to cope with an increasing number of local optima, there are several parameters which have an effect. Parameters such as the *number of sites* under investigation, the *number of bees* attributed to those sites, and the *differentiation* between sites and ‘elite’ sites are all factors which affect the searching behaviour of the algorithm to allow for greater flexibility as the modality of the problem landscape increases.

Results for BA are shown in Figure 2. Post-tuning, we find that the BA selects the *same* parameter configuration, regardless of the number of local optima present in the landscape. We then see that tuning has no effect on the ability of the algorithm to cope with increasing numbers of local optima. As long as the number of sites under investigation is greater than the number of optima, the algorithm is capable of dealing with modality. Coupled with the abandonment of ‘unpromising’ sites, this means that ‘too many’ sites are not detrimental to the exploration pattern of the algorithm.

We see the same pattern when increasing the ratio of local optima to the global optimum. As long as the number of sites under investigation covers the modality of the landscape, the BA is not hampered by increasing levels of attractiveness, regardless of parameter settings. The *patch size*

¹<http://dx.doi.org/10.6084/m9.figshare.696908>

Table 2: Mean (μ) and standard deviation (σ) of the exact average error of algorithm performance (both untuned (UT) and tuned (T)). Smaller values imply more robustness to changes in a specific characteristic.

		BFOA		Bees Algorithm		ES		GA		Harmony Search		PSO		SHC	
		UT	T	UT	T	UT	T	UT	T	UT	T	UT	T	UT	T
		μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
# of Local Optima	μ	0.118	0.003	0.001	8.8×10^{-6}	0.085	0.078	0.093	0.015	0.011	2.2×10^{-5}	0.025	0.014	0.266	0.072
	σ	0.011	0.001	2.1×10^{-4}	9.2×10^{-7}	0.028	0.026	0.033	0.008	0.005	2.9×10^{-5}	0.010	0.010	0.041	0.020
Dimensions	μ	0.754	0.417	0.216	0.073	0.542	0.544	0.420	0.529	0.364	0.263	0.420	0.157	0.577	0.589
	σ	0.388	0.360	0.202	0.069	0.345	0.346	0.233	0.401	0.271	0.204	0.307	0.145	0.261	0.371
Local Optima Ratio	μ	0.120	0.003	0.001	8.7×10^{-5}	0.084	0.082	0.079	0.007	0.007	0.002	0.025	0.016	0.284	0.088
	σ	0.021	0.002	2.3×10^{-4}	1.9×10^{-4}	0.012	0.012	0.006	0.006	0.003	0.004	0.004	0.011	0.045	0.011
Boundary Range	μ	0.317	0.022	0.001	0.001	0.097	0.093	0.125	0.021	0.048	0.001	0.076	0.022	0.446	0.305
	σ	0.213	0.033	1.3×10^{-4}	0.001	0.017	0.018	0.057	0.016	0.041	0.001	0.050	0.013	0.239	0.217
Smoothness	μ	0.260	0.010	0.004	0.001	0.110	0.102	0.154	0.021	0.018	0.001	0.043	0.014	0.349	0.112
	σ	0.089	0.005	0.002	4.3×10^{-4}	0.012	0.012	0.045	0.011	0.007	0.001	0.012	0.006	0.039	0.014

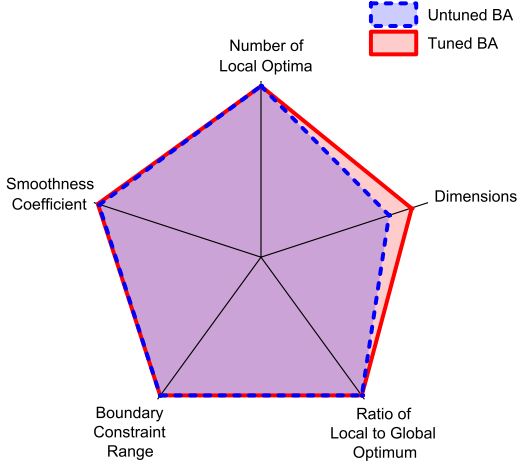


Figure 2: Summary of results for **bees algorithm**.

parameter of the BA controls the distance from a site bees are allowed to explore. This is the parameter which affects the search behaviour of the algorithm as boundary constraint size increases. The BA allows for full coverage of any sized search space, using *scout bees* to investigate new random sites to give ‘teleportation’ across the landscape. As with the number of local optima, we find the F-Races for the BA select the same parameter set for most boundary constraint sizes. We find that, post-tuning, the performance of the BA actually *decreases* slightly, suggesting the algorithm can cope less well with changes in boundary constraint size. We believe that the configurations may have become over-fitted to the landscapes used for tuning, and, while performance on the landscapes used for racing may have increased, the ability to search generalised landscapes has decreased.

Dimensionality provides the most significant result in terms of pre-tuning and post-tuning performance of the BA. We observe little change in performance at one to three dimensions - the point where the untuned algorithm is already performing well. As dimensionality increases beyond this, the effect of tuning becomes *increasingly* beneficial. We

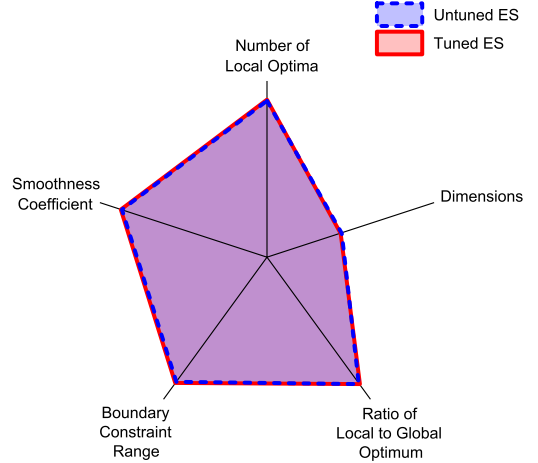


Figure 3: Summary of results for **evolution strategies**.

suggest that there is no increase in performance in other characteristics because these landscapes are *not challenging enough* to the BA to require adjusting the parameters. For the ranges of landscape characteristics on which we have tested the BA, it is clear that tuning generally makes little difference to the performance, as suggested by its original developer.

Evolution Strategies

ES has the smallest number of parameters of all the algorithms studied here (excepting the baseline algorithm, stochastic hill climbing). The two parameters this form of ES offers are (1) *population size* and (2) *number of children*. It is suggested (Cant-Paz, 2001) that altering these parameters adjusts *selection pressure* (that is to say, the *greediness* of the algorithm changes). The parameter configurations obtained through F-Racing are varied, implying that there do exist some configurations that are more successful than others. A range of configurations are selected across each characteristic - both in terms of different values for the two parameters, and different selection pressures when the two

parameters are combined.

Results for ES are shown in Figure 3. It is perhaps surprising to observe that the results of using the tuned parameters show little or no change in performance across all characteristics. There is a small decrease in average error as the number of local optima changes, but the standard deviation is similar for both untuned and tuned, suggesting that while the average error has decreased very slightly, the ability of the algorithm to cope with increasing numbers of local optima is unchanged. For all other characteristics post-tuning, there is little change in both average error and standard deviation across characteristics values (that is to say, the algorithm is no more capable of dealing with changes in these characteristics). This is perhaps consistent with the definition of the two parameters the algorithm offers - selection pressure can only affect the way in which ES explores local optima, and there is no control over the area that is explored around each point of interest, or any way to encourage the algorithm to rapidly explore an increasingly large search space.

We use a simple variant of ES, here, and there exist many other versions of the ES algorithm that offer a greater range of parameters (such as CMA-ES (Hansen and Kern, 2004)). ES clearly yields its best performance with an “out-of-the-box” parameter configuration, which means that it is quick to implement. However, our results suggest that there is little that can be done to improve the performance of this particular variant.

Genetic Algorithm

The performance of the GA increases post-tuning, coping significantly better with increasing numbers of local optima, increasing boundary constraint range and an increasing smoothness coefficient. Results for the GA are shown in Figure 4. The parameters of the GA are not as intuitively linked to the exploration pattern as many of the other algorithms in the study. This particular GA offers four configurable parameters: (1) *population size*, (2) *‘bits’ per parameter* in the representation, (3) *crossover rate* and (4) *mutation rate*. In experiments with a fixed number of objective function calculations, population size affects the number of generations the algorithm evaluates before terminating. A larger number of bits in a bit string representation allows more ‘precise’ solutions to be generated at the expense of having a representation which is less affected by mutation. Similarly to BFOA, there are a few configurations which re-occur across different characteristics and different characteristic values. It is probable that once a ‘good’ configuration has been found for a GA, it is applicable to ‘similar’ landscapes, which is consistent with the suggestion Goldberg (1989) that GAs are robust problem solvers, exhibiting approximately the same performance across a wide range of problems.

With increasing dimensionality, the GA initially shows

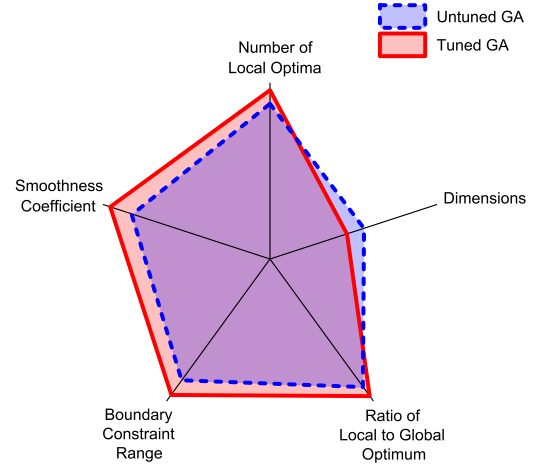


Figure 4: Summary of results for **genetic algorithm**.

promising results in terms of tuned performance, with a marked performance increase up to four dimensions. The benefit from tuning rapidly declines, however, until the tuned performance is *worse* than that of the untuned version. There are two possible explanations for this: the first is that the restriction on the number of objective calculations did not allow the F-Race algorithm to gather any meaningful performance data from the configurations. The second explanation is that we did not test a wide enough range of configurations - although two of the four parameters have definite ranges (mutation and crossover rates are percentages, thus generation was bounded between zero and one), so this is unlikely.

Harmony Search

The four parameters of HS all control different aspects of the search strategy. *Memory size* dictates how many promising solutions can be stored - effectively, how many potential sites of interest are retained by the algorithm. *Consideration rate* and *adjustment rate* control how new solutions are generated. The consideration rate is the percentage chance that a solution based on one in memory will be generated (conversely, *1-consideration rate* is the chance a random solution is generated instead). The adjustment rate is then the percentage chance that the randomly chosen solution from memory will be adjusted. If so, the fourth parameter, which controls the *maximum range* at which solutions can be adjusted, is used. If the adjustment does not occur, the considered solution potentially occupies an additional slot in the memory - thus increasing the chance that this solution may be chosen for consideration again. The interplay between these parameters is crucial, and it is somewhat hard to see how consideration rate and adjustment rate can directly affect the search strategy - unlike memory size and range,

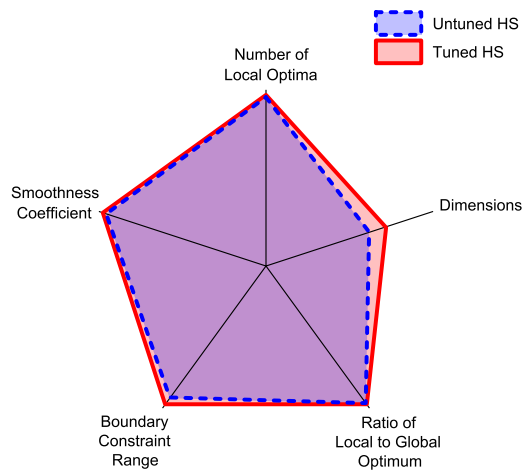


Figure 5: Summary of results for **harmony search**.

which are more obvious.

The results for HS are shown in Figure 5. HS, like the BA, offers some of the lowest ‘out of the box’ average error rates in this study. For most characteristics, there is little room for a performance increase post-tuning. Boundary constraint range proves to be the second-most challenging characteristic to HS pre-tuning, but post-tuning shows improved performance. The range values in all the configurations selected by F-Racing are much smaller than those in the ‘out of the box’ values, and this contributes significantly to the performance improvement when boundary constraint ranges are increasing. The consideration rate also decreases almost linearly as size increases - effectively, more random solutions are used instead of relying on the ‘memory’. These random solutions allow the solution pool to jump from one position in the search space to another, encouraging a wider search space, and explaining the significant improvement as boundary constraint range increases. Dimensionality also yields an improvement in the tuned parameter performance of HS, in terms of both average error and ability to cope, as it rises. High dimension problems (seven and above) have a much higher consideration rate than the successful configurations for lower dimensionality, suggesting that a focus on *exploitation* rather than *exploration* is beneficial to the HS when dimensionality is high. This is the opposite case of what happens with boundary constraint range, as discussed above.

Particle Swarm Optimisation

PSO in this form has four parameters; these control the *population size*, the *maximum velocity* of a particle, the bias towards the *particle best solution* and the bias towards the *global best solution*. With these parameters, it is possible to control the coverage of a search space (the number of parti-

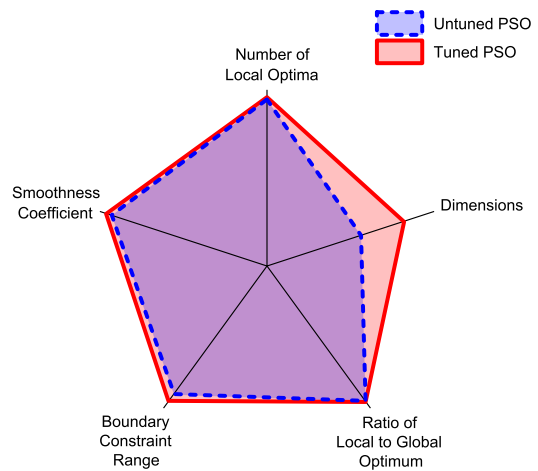


Figure 6: Summary of results for **particle swarm optimisation**.

cles), enforce a large search area of a small search area for each particle (the maximum velocity), and, through manipulation of the local and global best solution bias, control the capability of the algorithm to converge on a single solution or explore several areas of interest (optima avoidance).

Results for PSO are shown in Figure 6. The parameters used cover a broad range of search behaviours, and, as such, we would expect to see a large improvement in particle swarm performance post-tuning. This holds true for most of our characteristics. Results for the number of local optima, for example, show a reasonable decrease in average error as the number of local optima increases, yet the standard deviation demonstrates no change, indicating that the algorithm is no more capable of dealing with increasing numbers of local optima post-tuning. Performance of PSO greatly improves on dimensionality post-tuning, in terms of both average error and ability to cope as it grows. The F-Race algorithm for PSO selects the same configuration for all values of dimensionality (except for 2 dimensions), implying that there is no specific parameter that requires adjustment to cope with the increase in dimensionality, but selecting a configuration which provides good *exploration* allows PSO to perform well as the size of the search space increases exponentially.

This trend continues across all characteristics, with F-Races often selecting the same configurations, regardless of characteristic values. As with the other swarming algorithms, we suggest that once a good configuration has been found, it is able to deal with a wide range of problems of a similar nature, regardless of the specific characteristics. The configurations selected are all varied in their parameters, and it is unexpected to see that there is no pattern to maximum velocity as boundary constraint range increases. This

is possibly because maximum velocity is an upper bound, and there are particles with randomly generated velocities below the maximum, so this parameter is less significant than it may initially appear. It would perhaps be interesting to consider the effect of having a *minimum* velocity on the increase in boundary constraint range, although this would also severely hamper exploitation.

Stochastic Hill-Climbing

With only a single parameter - the *range* at which new solutions are generated - the SHC algorithm does not offer a large amount of customisation. This single parameter is directly linked to the search pattern and nothing else, and as there are no other parameters there is no interplay between parameters to consider. Arguably, therefore, SHC should prove the easiest algorithm to tune.

Results for SHC are shown in Figure 7. All characteristics, barring dimensionality, show an improvement post-tuning. As the neighbourhood size is the range at which new solutions are generated, it is unsurprising that tuning improves algorithm performance as boundary constraint ranges change. As the number of objective function calculations is limited, despite having a larger neighbourhood size, the ability of the algorithm to effectively explore larger environments is still restricted, therefore the average error does not decrease by as much as may be expected, and the ability of the algorithm to deal with increasing search space sizes improves only slightly.

SHC demonstrates a large increase in performance and a greater ability to cope with more optima (a reduced standard deviation) post-tuning. The parameter configurations selected for the number of local optima, the ratio of local optima *and* the smoothness all have a neighbourhood size of around 50% of the search space size. We suggest that the performance improvement for all of these characteristics is actually derived from the algorithm having configured itself properly for the search space size used as a default for all other characteristics, rather than tuning itself to best perform on any specific characteristic.

Conclusions and Future Work

In this paper we have built on previous studies of the performance of nature-inspired algorithms on fitness landscapes with different characteristics. Earlier work explored ‘out of the box’ parameter configurations, and we further develop this by using an automated parameter configuration methodology. This allows us to study the effect of tuning on different algorithms, contributing significantly to the debate on when and how it is beneficial to tune specific algorithms.

We observe that algorithms broadly fall into three categories: (1) algorithms which do not benefit from tuning (ES), (2) algorithms which benefit from tuning when the problem is “difficult” enough to require tuning (i.e. when

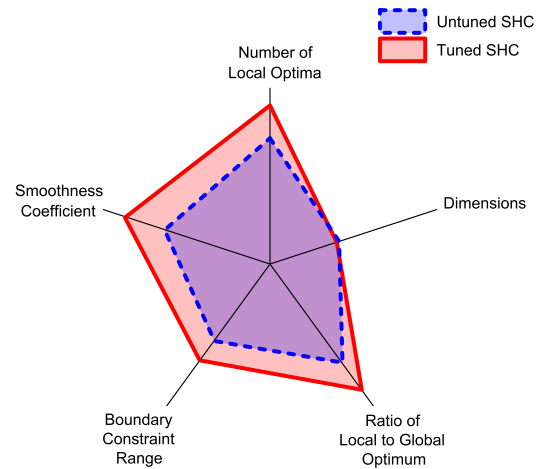


Figure 7: Summary of results for **stochastic hill-climbing algorithm**.

untuned performance is poor) (BA, HS, PSO), and (3) algorithms which always benefit from tuning (BFOA, GA, SHC). Dimensionality often offers the most significant improvement post-tuning in algorithms, particularly those with parameters that increase the *breadth* of search space (swarming algorithms are significantly better here than evolutionary algorithms).

The methodology presented here is easy to implement, is computationally inexpensive, and offers considerably more information on the performance of an algorithm than using a standard set of benchmark problems. We hope that it will offer a foundation for future work on the experimental comparison of nature-inspired algorithms, as well as a useful set of heuristics for practitioners to use in order to decide when and how to tune their methods.

References

- Adenso-Diaz, B. and Laguna, M. (2006). Fine-tuning of algorithms using fractional experimental designs and local search. *Operations Research*, 54(1):99–114.
- Akay, B. and Karaboga, D. (2009). Parameter tuning for the artificial bee colony algorithm. In Nguyen, N., Kowalczyk, R., and Chen, S.-M., editors, *Computational Collective Intelligence. Semantic Web, Social Networks and Multiagent Systems*, volume 5796 of *Lecture Notes in Computer Science*, pages 608–619. Springer Berlin Heidelberg.
- Bäck, T. and Schwefel, H.-P. (1993). An Overview of Evolutionary Algorithms for Parameter Optimization. *Evolutionary Computation*, 1(1):1–23.
- Barr, R., Golden, B., and Kelly, J. (1995). Designing and reporting on computational experiments with heuristic methods. *Journal of Heuristics*, 1:9–32.
- Beyer, H.-g. and Schwefel, H.-p. (2002). Evolution strategies. *Natural Computing*, 1:3–52.

- Birattari, M. (2009). *Tuning metaheuristics: a machine learning perspective*. Springer.
- Birattari, M., Stützle, T., Paquete, L., and Varrentrapp, K. (2002). A racing algorithm for configuring metaheuristics. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '02*, pages 11–18, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Birattari, M., Yuan, Z., Balaprakash, P., and Stützle, T. (2010). F-race and iterated f-race: An overview. *Experimental methods for the analysis of optimization algorithms*, pages 311–336.
- Brabazon, A. and O'Neill, M. (2006). *Biologically inspired algorithms for financial modelling*. Springer-Verlag.
- Brownlee, J. (2011). *Clever Algorithms: Nature-Inspired Programming Recipes*. Lulu.
- Cant-Paz, E. (2001). Migration policies, selection pressure, and parallel evolutionary algorithms. *Journal of Heuristics*, 7(4):311–334.
- Crossley, M., Nisbet, A., and Amos, M. (2013). Fitness landscape-based characterisation of nature-inspired algorithms. In Tomassini, M., Antonioni, A., Daolio, F., and Buesser, P., editors, *Proceedings of the 11th International Conference on Adaptive and Natural Computing Algorithms (ICANNGA'13)*, Lausanne, Switzerland, April 4-6, 2013. *Lecture Notes in Computer Science*, Vol. 7824, pages 110–119. Springer.
- Eiben, A. E. and Smit, S. K. (2011). Parameter tuning for configuring and analyzing evolutionary algorithms. *Swarm and Evolutionary Computation*, 1(1):19–31.
- Gallagher, M. and Yuan, B. (2006). A general-purpose tunable landscape generator. *IEEE Transactions on Evolutionary Computation*, 10(5):590–603.
- Geem, Z. and Kim, J. (2001). A new heuristic optimization algorithm: harmony search. *Simulation*, 76(2):60–68.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley.
- Hansen, N. and Kern, S. (2004). Evaluating the cma evolution strategy on multimodal test functions. In *Parallel Problem Solving from Nature-PPSN VIII*, pages 282–291. Springer.
- Hendtlass, T. (2009). Particle swarm optimisation and high dimensional problem spaces. In *2009 IEEE Congress on Evolutionary Computation, CEC'09*, pages 1988–1994. IEEE.
- Horn, J. and Goldberg, D. (1994). Genetic algorithm difficulty and the modality of fitness landscapes. In *Foundations of Genetic Algorithms 3*.
- Jones, T. and Forrest, S. (1995). Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 184–192.
- Kennedy, J. and Eberhart, R. (1995). Particle swarm optimization. In *Neural Networks, 1995. Proceedings. ...*, pages 1942–1948.
- Koster, C. and Beney, J. (2007). On the importance of parameter tuning in text categorization. *Perspectives of Systems Informatics*, pages 270–283.
- Kukkonen, S. and Lampinen, J. (2005). GDE3: The third Evolution Step of Generalized Differential Evolution. *2005 IEEE Congress on Evolutionary Computation*, pages 443–450.
- Leung, F. H., Lam, H., Ling, S., and Tam, P. K. (2003). Tuning of the structure and parameters of a neural network using an improved genetic algorithm. *Neural Networks, IEEE Transactions on*, 14(1):79–88.
- Lobo, F. G., Lima, C. F., and Michalewicz, Z. (2007). *Parameter setting in evolutionary algorithms*. Springer Verlag.
- Malan, K. M. and Engelbrecht, A. P. (2009). Quantifying ruggedness of continuous landscapes using entropy. In *2009 IEEE Congress on Evolutionary Computation*, pages 1440–1447. IEEE.
- Maron, O. and Moore, A. (1993). Hoeffding races: Accelerating model selection search for classification and function approximation. *Robotics Institute*, page 263.
- Maron, O. and Moore, A. W. (1997). The racing algorithm: Model selection for lazy learners. *Artificial Intelligence Review*, 11:193–225.
- Merz, P. (2000). Fitness landscape analysis and memetic algorithms for the quadratic assignment problem. *Evolutionary Computation, IEEE*, 4(4):337–352.
- Morgan, R. and Gallagher, M. (2010). When does dependency modelling help? Using a randomized landscape generator to compare algorithms in terms of problem structure. In et al Schaefer, R., editor, *PPSN XI*, pages 94–103. Springer-Verlag.
- Nannen, V., Smit, S. K., and Eiben, A. E. (2008). Costs and benefits of tuning parameters of evolutionary algorithms. In *Parallel Problem Solving from Nature-PPSN X*, pages 528–538. Springer.
- Passino, K. (2002). Biomimicry of bacterial foraging for distributed optimization and control. *IEEE Control Systems Magazine*, 22(3):52–67.
- Pham, D., Ghanbarzadeh, A., and Koc, E. (2006). The Bees Algorithm: A Novel Tool for Complex Optimisation Problems. In Pham, D., Eldukhri, E., and Soroka, A., editors, *Intelligent Production Machines and Systems*, pages 454–459.
- Ridge, E. and Kudenko, D. (2010). Tuning an algorithm using design of experiments. In *Experimental methods for the analysis of optimization algorithms*, pages 265–286. Springer.
- Smit, S. K. and Eiben, A. E. (2009). Comparing parameter tuning methods for evolutionary algorithms. In *Evolutionary Computation, 2009. CEC'09. IEEE Congress on*, pages 399–406. IEEE.
- Yuan, B. and Gallagher, M. (2004). Statistical racing techniques for improved empirical evaluation of evolutionary algorithms. In *Parallel Problem Solving from Nature-PPSN VIII*, pages 172–181. Springer.